

SAN Value Model

ITCentrix, Inc.

Infolink Overview

Table of Contents

Introduction	2
Objectives	2
Prerequisites	2
Infolink Overview	3
Module I: What Is Value Selling?	3
Module II: The Five Stages of Value Selling	3
Module III: Using the Value Software Tool	3
Module IV: Case Studies	4
Naming Conventions and Icons	4
Icon/Name	4
Description	4
What Is Value Selling?	5
How the Perception of Value is Changing	6
Catalysts for Change:	6
Practical Examples:	7
Purpose	8
What's in this Module?	8
Stage I: Establish Interest and Intent	9
Stage II: Engage In Discussion of IT Value	9
Stage III: Advise Customer on New Approach	11
What is the ITcentrix Value Contribution Model?	11
How is the Model Applied to a Compaq SAN?	12
Stage IV: Apply the Model to a Customer Environment	14
Typical Customer Storage Environments	14
How SAN Addresses the Problems	16
Business Impacts:	16
Value of SAN by Functional Improvements	17

Stage V: Agree on an Action Plan	18
Action Plan List	18
Purpose	19
What's in this Module?	19
A Conceptual View of the Software Tool	19
Step 1: Assessing IT Infrastructure	21
Step 2: Assessing Operational Cost	22
Step 3: Assessing the Target SAN Infrastructure	26
Step 4: Assessing Application Value	26
Application Value Multiplier	29
EXAMPLE – Typical Customer Interaction	29
Step 5: Applying Application Value to Availability	31
Step 6: Applying Application Value to Flexibility	33
Step 7: Delivering Results	34
Purpose	35
What's in this Module?	35
Assumptions	36
Glossary Terms	39
Business Value Contributors –	39
Operational Costs	39
Service Levels	39
Business Flexibility	39
Distributed Storage	39
Collocated Storage	39
Homogeneous SAN (San of Today)	39
Limited Heterogeneous SAN (SAN of 2000)	39
Future SAN (San of 2002)	40
Storage Area Network (SAN)	40

Introduction

Current IT industry analyst data supports the premise that in order to penetrate high-level enterprise customer accounts, IT vendors must be able to sell directly to the highest levels of management, in terms of "business value."

In response to this need, Enterprise Storage Marketing has invested in the development of a "business value" selling tool. The tool was designed to equip Compaq's sales reps and storage specialists with a means for identifying, soliciting, and quantifying the "business value" of a storage architecture and strategies within a customer's enterprise computing environment.

Objectives

After a full review of these infolinks, users who meet the prerequisites will be able to:

- 1.) Define the term "value selling" in the IT business framework.
- 2.) Identify and define the five stages of value selling.
- 3.) Identify the key components of the SAN Value Model.
- 4.) Apply the model to target Compaq SAN infrastructures.
- 5.) Calculate application values given specific customer case studies

Prerequisites

These infolinks are designed for Compaq storage specialists and sales representatives who have at least two years of success at selling Compaq StorageWorks Solutions. Users should also have specific product training in Compaq SAN solutions.

Infolink Overview

Module I: What Is Value Selling?

What Is Value Selling takes a close look at the relationship between IT business value and value selling and how value selling directly connects to the business and measures change.

Module II: The Five Stages of Value Selling

The Five Stages of Value Selling includes a detailed description of each stage as well as tips and suggestions for each.

- Stage 1: **Establish** – an interest and need
- Stage 2: **Engage** – in a substantive discussion of value
- Stage 3: **Advise** – your client on a new approach
- Stage 4: **Apply** – the approach to your customer's environment
- Stage 5: **Agree** – on a follow-up action plan

Module III: Using the Value Software Tool

The following seven steps are used to assess a customer's storage infrastructure:




Step 1: Assessing the overall IT Infrastructure
Step 2: Assessing Operational Costs
Step 3: Assessing the Target SAN Infrastructure
Step 4: Assessing the Value of SAN Applications
Step 5: Applying Application Value to Availability
Step 6: Applying Application Value to Flexibility
Step 7: Delivering Results

Module IV: Case Studies

This section is composed of case studies based on real information from actual client data.

Naming Conventions and Icons

The purpose of this section is to provide the user all the naming conventions and icons within the Infolinks. Refer to the table below.

Icon/Name	Description
Note:	Indicates that the information that follows is very important.
Reference:	A reference pointer refers to additional information within the guide
	This icon represents a <i>tip or suggestion</i> .
	This icon represents a sample screen from the tool.
	This icon represents a notes page.

Introduction to Value Selling

Module 1

What Is Value Selling?

Value Selling is the process of helping the customer connect IT to the business and measuring the change. The process involves considering the customers' IT in business value terms and helping the customer to think logically about ways to understand the link between business and IT. While ultimately the goal is to sell a solution, value selling demands that you establish credibility with the customer by providing advice and counseling on new ways to think about IT planning; independent of products and services.

Why Value Selling?

The reasons for value selling are generally accepted. By selling value, you can appeal to higher levels within the organization, command better prices and establish better relations with the customer. Typically, no one person has an overall view of an enterprise, and value selling can open the door to many functions within an organization. Examples of this are:


- Applications Groups
- CIOs
- Business Management



Presenting this comprehensive view strengthens your credibility and adds value to the customer

How the Perception of Value is Changing

Through case studies with dozens of Global 2000 (and smaller) corporations and other government and non-profit institutions, it is becoming increasingly clear that the perception of IT Value is changing (see figure below).



Source: ITcentrix, Inc. - 1999

Catalysts for Change:

- Decades of IT investment focused on Internal Issues
- The “Web effect” changes priorities from internal to external
- Customers are re-architecting IT to face the customer
- Customer Service Levels and Time-to-market are becoming more important

Practical Examples:

- What would it have meant to barnesandnoble.com to get its Web site up six weeks earlier?
- What’s more important to eBay, cost savings or availability?
- Do you have a Web example du jour?



Re-enforce to customers that operational cost savings are still crucial. Emphasize, however that other value factors must be considered and that traditional value tools like ROI and TCO, while complimentary, only address part of the equation. Customers today need to avoid “rear view mirror” planning and anticipate changes to their business models.

The Five Stages of Value Selling

Module 2

Purpose

The purpose of this module is to provide storage specialist and sales representatives with a detailed understanding of the five stages of value selling, as well as guidelines for value selling to customers.

What's in this Module?

The following stages are included in this module.

Topic	See Page
Stage 1: Establish Interest and Intent	
Stage 2: Engage In Discussion of IT Value	
Stage 3: Advise Customer On New Approach	
Stage 4: Apply The Model to a Customer Environment	
Stage 5: Agree A Plan of Action	

Stage I: Establish Interest and Intent

Establish with customers that you would like to explore the "business value of IT."
Start by explaining to customer:

- That you have worked with an independent firm, ITcentrix Inc. to develop a strategic model to assess the value of a storage infrastructure.
- That you would like to discuss various aspects of the customer's IT infrastructure with the intent of assessing the business value contribution of a storage approach.
- That the model boils value down to three factors:
 - Operational Cost
 - Service Levels
 - Business Flexibility



Help the customer understand that this model was developed from interactions with numerous large and medium-sized companies and that the exercise has value in and of itself.

Stage II: Engage In Discussion of IT Value

During this stage, engage the customer in an initial discussion of value by asking a series of questions.

- Does your company formally measure IT value?
- How does your company measure IT value?
- What tools do you use ? TCO? ROI? Others?
- Do you formally track the value of an IT project? How?
- Is the perception of IT value changing within your company?
- Is your company increasingly willing to invest in IT to increase revenue or... customer satisfaction?
- Do your competitors look at IT value differently than your company? How?

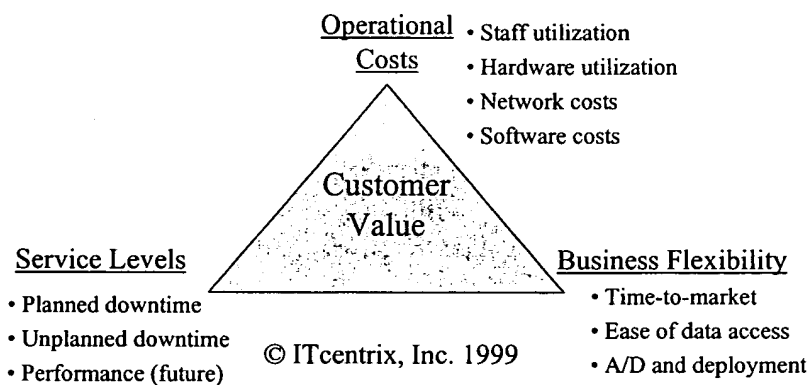
Reference: See Appendix E - Science of Selling-Opportunity Assessment Questions

Stage III: Advise Customer on New Approach

What is the ITcentrix Value Contribution Model?

The ITcentrix Value Model is a tested methodology that combines cost, technology and business modeling to more accurately predict how changes in IT will impact business results. It is a tool developed by an independent company built within a strategic framework to assess IT's contribution to business value. Data from the model is calibrated using real world studies of Global 2000 and smaller corporations.

The ITcentrix Value Contribution Model



The model uses assumptions based on actual customer data and allows users to make changes to reflect specific applications and environments (e.g. application value, levels of availability, planned downtime, network cost, etc.).

The following table details the three high-level Business Value Contributors and how they are quantified and analyzed to assess the value contribution of a particular technology approach.

Business Value Contributors	Metric	Key Ingredients
Operational Costs	Fully loaded cost savings	Staffing and storage capacity utilization
Service Levels	Higher revenue or productivity from improved availability	Availability at the end user level
Business Flexibility	Higher revenue or productivity from faster application deployment	Rate of change to application environment

Each of these high-level value contributors contains numerous sub-elements and data points solicited from actual customer situations. These factors are assessed to develop an accurate view of current technology approaches (*The Base Case*) and compared to alternatives. The explicit intent of the model is to allow customers to compare tradeoffs of changes to the Base Case in business value terms. [All components of the model are quantified in value terms and represent real dollars - e.g. cost savings, revenue potential and/or productivity gains].



Keep it simple for customers. Boil value down to these three factors. Remember the key ingredients of each Value Contributor and focus the discussion on these factors.

How is the Model Applied to a Compaq SAN?

The tool quantifies and assesses the value contribution of the following storage topologies:

- ***Distributed*** – Physically de-centralized (outside of a data center), servers managed centrally, storage directly attached to servers and managed locally, backup over the corporate network.
- ***Collocated*** – Physically centralized (inside a data center), servers managed centrally, storage directly attached to servers and managed centrally (by IT Professionals), backup over a local network within the data center.
- ***Homogeneous SAN (SAN of Today)*** – Shared, pooled storage attached to multiple Compaq servers running a homogeneous operating system using Compaq storage. Any-to-any connectivity between servers and I/O (disk and tape) using a common serial switched fabric and network switches; static allocation of server and I/O resources.
- ***Limited Heterogeneous SAN (SAN of 2000)*** – Shared, pooled storage attached to multiple Compaq and non-Compaq servers running Unix (some flavors) and/or NT operating systems using Compaq storage. Any-to-any connectivity between servers and I/O (disk and tape) using a common serial switched fabric and network switches; static allocation of server and I/O resources with improved management function.
- ***Future SAN (SAN of 2002)*** - Shared, pooled storage attached to multiple Compaq and non-Compaq servers running multiple operating systems using multiple vendors' storage devices. Single point of control for dynamic allocation of server and I/O resources to optimize service levels and availability. Unified SAN OS or Global OS support.



Remember that most customers are going to have some mix of distributed and collocated storage. Typically IT Professionals have a much better handle on the storage that is inside the data center (i.e. collocated) and much of the discussion will center on this class of storage. A Compaq SAN solution will fare quite well (value-wise) relative to Distributed and Collocated storage so these will be your "Sweet Spots."

Customers will often also have some EMC-like storage connected to their Open Systems. This is often referred to as "Consolidated." The value proposition of a Compaq SAN (of today) has many similarities to an EMC (or Hitachi) consolidated approach (e.g. the benefits of shared storage pools across more than two servers) and you should position SAN as "the great equalizer." While there are many differences in the two approaches, the most important point is that **you no longer have to search exclusively for distributed or collocated opportunities. You can now effectively compete for pooled storage deals with EMC and Hitachi.**

Stage IV: Apply the Model to a Customer Environment

Typical Customer Storage Environments

In an ITcentrix survey completed in September 1999, we asked large and medium-sized customers "What is your single biggest storage problem?" The following table shows the responses:

Q42a. What is your biggest storage problem today?

	Frequency	Percent	Valid Percent
Backup window problems	79	26.3	26.3
Difficulty managing growth	66	22.0	22.0
Performance	24	8.0	8.0
Load balancing is difficult	14	4.7	4.7
Data availability	5	1.7	1.7
Lack of an enterprise storage strategy	4	1.3	1.3
Reconfiguration is difficult	4	1.3	1.3
Scalability	4	1.3	1.3
Recovery	3	1.0	1.0
Can't find or retain staff	2	.7	.7
Service levels	2	.7	.7
Storage is underutilized	1	.3	.3
Security	1	.3	.3
Other 1	91	30.3	30.3
Total	300	100.0	100.0

Note: Other includes "no problems" and items that comprise less than 5% of the total responses each.

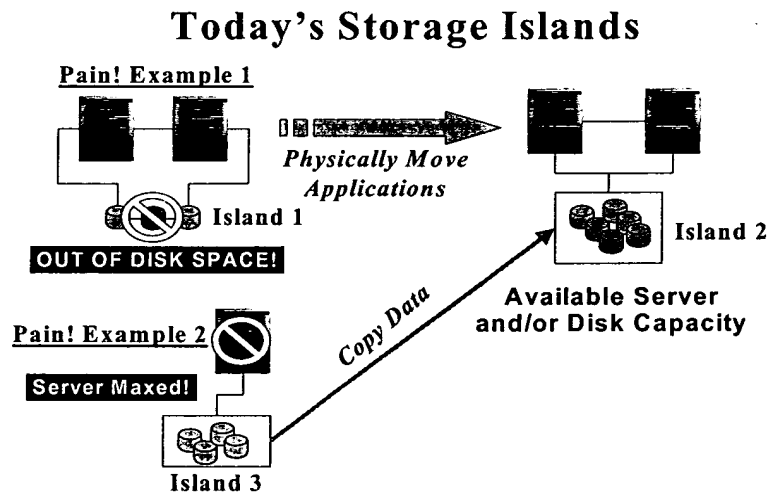
The survey shows that by far the two biggest customer issues are:

- Backup Window
- Difficulty Managing Growth.

The backup window problem is fairly straightforward in that customers frequently have many backup procedures for different types of server and storage platforms. This creates inefficiencies and elongates backup times. The goal of many customers is to reduce the number of procedures and increase the commonality in backup approaches.

Backup window problems have fairly significant business implications. Consider that backup window problems frequently cause so-called "Batch" overruns (i.e. system maintenance creeping into business operation hours). Batch overruns can cause unplanned downtime and cost hard dollars. In turn, many companies will forego a backup if it eats into normal business operations. This approach exposes a company to increased risk.

“Difficulty managing growth” is a bit more complicated. The following example demonstrates typical customer issues with managing growth:



The example is drawn from real world case studies and involves collocated storage (separate storage islands) in two scenarios:

- 1.) An out of disk space problem.
- 2.) The server is running at or above acceptable threshold levels.

In each case, customers have three unappealing choices:


1. Buy more disk capacity
2. Physically move the applications (or re-cable the systems) to a location with additional capacity
3. Copy data over an already clogged corporate network to a location with additional capacity

Either the customer is buying storage when free space exists somewhere on the floor or the customer is forced to re-configure systems (at 3AM on July 4th). In either case, inefficiencies mount.

How SAN Addresses the Problems

SAN directly solves these issues cited above. In the case of backup window problems, by using a pooled storage approach. Common software and backup procedures can be implemented thereby dramatically simplifying the backup process and reducing pressures on backup windows.

As it relates to managing growth, consider the following depiction of an any-to-any topology:



In the example, if a disk or server reaches capacity limits, the SAN topology allows the system to redirect access to data (today manually and in the future dynamically).

Business Impacts:


1. Simplified Operations
 - A. Shared pooling
 - B. Better capacity utilization
2. Improved Availability
 - A. Balanced systems
 - B. Fewer room for mistakes
 - C. Enables automation (future)
3. Faster Time-to-market
 - A. Direct access to data
 - B. Easier application development and deployment
 - C. Bronchitis

Think about what a customer with 50 or 100 or 200+ servers is facing. Typically, most of the servers will have attached storage that is not accessible by any other server. The above examples of “out of disk space” or “server is maxed out” probably occur quite frequently. While SANs don’t totally eliminate these problems (because the customer will still be creating so-called “Cloudlets” with multiple SANs) they are a major step toward a single logical view of storage” and will have major business impacts.

Value of SAN by Functional Improvements

The ITcentrix Value Model quantifies the benefits of an any-to-any (any I/O to any server) SAN topology. Specifically, through extensive case studies with major corporations we assess the value of a SAN infrastructure and its “functional timeline.” By functional timeline, we refer to the following Compaq SAN software roadmap:

Compaq SAN Timeline

1999	2000	2001	Future 
✓ Sequential Vol Sharing (EVM)	✓ Improved Static Vol Mgmt (EVM)	✓ SAN-Wide Virtualization	✓ SAN Fabric Mgmt (basic)
✓ Static Vol Mgmt (SWCC/SSP)	✓ Copy Services (EVM/DRM)	✓ Global LVM (moving volumes in a SAN)	✓ SAN Security (basic)
✓ SAN Fabric Mgmt (basic) (SWCC)	✓ Snapshot & Cloning	✓ Dynamic Vol Mgmt	✓ Dynamic W/L Mgr
✓ SAN Security (basic) (SSP)	✓ Impved Vol Sharing	✓ Workload Mgr	✓ Dynamic Allocation
✓ Base Copy Services (DRM)	✓ Alternate Pathing (SecurePath)	✓ Global SAN File Mgmt	
	✓ Improved Admin	✓ Global File System	
	✓ Improved Security	✓ Static SAN SRM	
	✓ Improved Fabric Mgmt	✓ Read only access to Heterogeneous file types	
	✓ Simpler Switch Mgmt		
	✓ Cascading		
	✓ Multi-path thru Switches		
	✓ Global LVM (reallocate, allocate vol in a SAN)		
	✓ Initial SAN Storage Resource Mgmt		

Stage V: Agree on an Action Plan

Action Plan List

Gain agreement to apply the model to the customer IT infrastructure and present an analysis of the results. This analysis can be included as a “value section” in a formal proposal. The Action Plan should be designed to:

- Engage customer organizations in discussion of SAN value.
- Support customer decision-making.
- Expose tradeoffs of current storage infrastructure and against alternatives.
- Examine the benefits of a Compaq SAN in business value terms.
- Assess SAN applicability to the specific customer’s environment.
- Provide direct inputs to customer business cases.

Using the Value Software Tool: Six Steps

Module 3

Purpose

The purpose of this module is to provide the storage specialists and sales representatives with an overview of the software tool and the six steps necessary for its effective use.

What's in this Module?

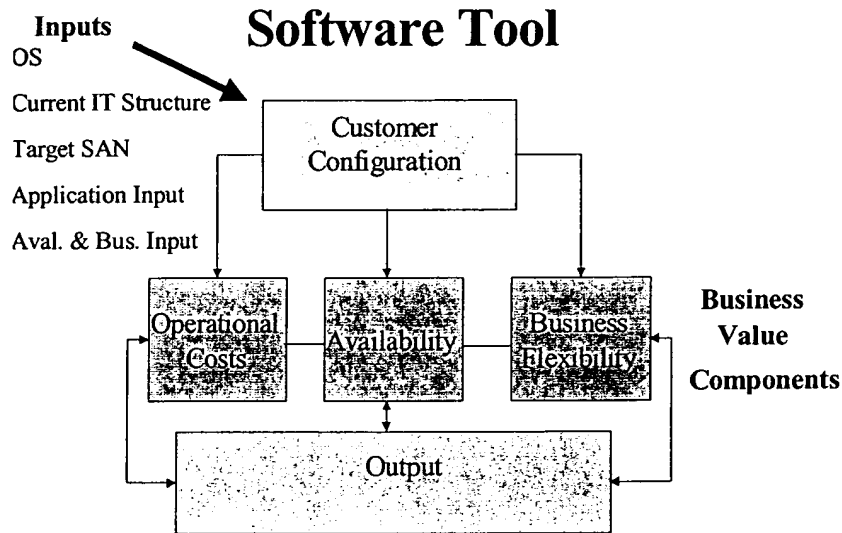
The following sections are included in this module.

Topic	See Page
A Conceptual View of the Software Tool	
Step 1: Assessing IT Infrastructure	
Step 2: Assessing Operational Costs	
Step 3: Assessing the Target SAN Infrastructure	
Step 4: Assessing the Value of SAN Applications	
Step 5: Applying Application Value to Availability	
Step 6: Applying Application Value to Flexibility	
Step 7: Delivering Results	

A Conceptual View of the Software Tool

The Software Tool works by inputting specific data, which generates a value associated with the three main components of the model; Operational Cost, Service Levels (Availability) and Business Flexibility based on a given customer scenario.

The following diagrams provide a structural view of the software tool.



The Software Tool uses inputs from seven screens in the tool including:

- Company Information – Includes basic background reference points
- Current IT Infrastructure – Establishes current server, OS and storage installation
- Staff Input – Helps assess the storage capacity managed per person
- Target SAN Infrastructure – Identifies the SAN opportunity and scale
- Application Input – Focuses on the application(s) environment
- Calculating Application Value
- Availability and Business Flexibility – Applying the customer's application value to derive a customized value for higher availability and business flexibility

In addition, the software contains seven other screens; two summary screens and five screens of outputs. These screens can be used in proposals as appropriate and include:

- Input Summary: A summary of the input provided in the previous screens
- Assumptions - that the model uses that are not changeable by the user
- Compaq SAN Value \$: A chart showing the four year absolute dollar value of a Compaq SAN broken down by storage topology (Collocated, SAN of Today, etc.) and value component (Cost, Availability, Flexibility).
- Compaq SAN Value %: A chart showing the percentage dollar value of a Compaq SAN broken down by storage topology (Collocated, SAN of Today, etc.) and value component (Cost, Availability, Flexibility).
- Staff Efficiency by Environment: A chart showing the customer's disk capacity managed per person – GB Managed Per Person (today) compared to projections if using a Compaq SAN.

- Availability Assumptions: A chart showing the availability assumptions used in the model. The chart shows: 1) Availability during normal business operations and 2) Scheduled downtime. Availability figures shown are end-user availability data.
- Components of Cost: A chart detailing the absolute four year dollar costs for the major factors of cost including: Server and software, disk and storage software, tape device, backup network, switch, server staff and storage staff costs. All costs are fully loaded and depreciated on a four-year schedule.

What follows is a description of the major steps involved in using the software. Each step may involve more than one screen in the tool.

Step 1: Assessing IT Infrastructure

Chances are that in each sales situation this is information you're already capturing. The step presumes you are interfacing with an individual that is an IT Professional with substantial knowledge of the customer's overall business, its IT installation, applications, storage specifics and overall IT direction. The step involves two input screens including:

- Screen 1: Basic company background information (revenue, # of employees, etc.).
- Screen 2: Information about the company's current IT infrastructure; with an emphasis on the types of servers and storage attached.

For the purposes of populating the tool, use the following parameters:

- Focus on storage that is under the management of IT Professionals.
- While the model doesn't exclude any platform...focus on the NT and Unix opportunity. Note: The model currently requires a separate analysis for the customer's Unix and NT storage infrastructures.
- Pick a Data Center (or centers) with which the individual has "adequate" familiarity.
- Expect that no one individual is going to have all the answers.
- Expect to get agreement to follow up with other individuals in the organization to "fill in the blanks."



Remember that The Model is as much a process as it is a "value generator." Your goal is to begin to develop a credible view of the

organization. If possible (and if time permits) you can capture input from multiple constituencies to strengthen your knowledge and business case. The process and model give you good justification to interface with Application Development Professionals, Line-of-business Managers and Senior IT Managers.

The typical approach to assessing a customer's IT infrastructure is to ask a series of questions that normally any customer can answer (at least generally) such as:

- 1.) Can you start by telling me about your IT infrastructure?
- 2.) How many servers and what types of servers (vendors and OS) do you have?
- 3.) Are the servers all in a centralized location?
 - A.) How many centralized locations and types of servers in each?
- 4.) How is the storage attached to these servers?
 - A.) Server specific (i.e. non-pooled)
 - B.) Pooled or Consolidated (i.e. multiple servers connected to a group of storage)
- 5.) How much storage (usable GB) is in each "group" (where a group is a set of storage attached to some N # of servers).

If the customer doesn't know the storage capacity...suggest that typically you see Unix servers with an average of **80-100GB** per system and NT servers with an average of **20-30GB** each. Ask the customer if this sounds about right for their situation.

Step 2: Assessing Operational Cost

This section is designed to ascertain details relevant to the Operational Costs Module. The module closely approximates a Total Cost of Ownership (TCO) analysis by including staff costs. The module is designed to be able to accept inputs from other TCO models and should be used in conjunction with such tools.

Start this section by explaining to the customer that the model uses inputs from the previous set of questions and makes certain assumptions about the customer's use of IT. Stress that any input of the model can be changed to reflect the customer's specific situation, but in the interest of time it might be wise to simply use some of the "defaults" for the first pass. You can always go back and further refine the model's inputs.

By far, the most important aspect of the Operational Cost Assessment is STAFFING. We've found this is the area that the customer typically underestimates in terms of cost impacts. By helping the customer uncover hidden staffing costs you can add significant perspective and value.

Inputs 3a and 3b on the STAFF INPUT screen are used to calculate your customer's GB's managed per person data (shown in the box on screen). You can compare this figure to industry averages by **REFERENCING THE CHART BELOW**.

Start by asking a series of questions related to staff including:

1.) For the Servers and Storage that we discussed previously, how many IT staff members (get full time equivalents) do you have looking after those resources? What is the total IT Staff?

Help the customer sort out the types of staff with the following potential "hit list."

- * System Administrators
- * Operations
- * Database Administrators
- * Applications Professionals
 - Applications Support
 - Applications Development
 - Applications Deployment
- * Management

2.) What percent of their time is spent doing Storage-related activities?

Nine out of ten times the answer will be "not much" or "**less than 10%.**" Your job is to "tease out" the real costs here. If done properly, the customer will often acknowledge that **25 – 55%** of staff time is spent managing storage. To get the customer thinking about hidden costs, go back to the above list and begin to assign storage-related tasks as follows:

- * Backup (operators) * Data Placement (DBAs) * Recovery (Apps)
- * Restore (systems) * Implementation (systems) * Data Integrity (Apps)
- * Tuning (systems/DBAs)* Administration (systems)* Finding Data (Apps)
- * Design (systems) * Acquisition (management) * Finding Capacity (Apps)
- * Security (DBAs) * Disaster recovery (systems)

Note: These tasks will never be a perfect match with these titles so be flexible. Make sure Management Time is included as applied to the amount of time managers spend managing storage professionals.

During this process, prompt the customer with questions such as:

- Who handles backups?
- What about restore (the most important part of a backup)?
- What happens when you run out of disk space?
- How often do you have to recover from a failure or error?
- When you add an application, is it hard to find available server and storage capacity?
- Who handles this task?
- Is it common that when you add an application you never know how much storage you need until the application is ready to be deployed?



Write down the % of time spent on each task and add up the totals for each job function. Keep good enough track to refer to your notes after the meeting as this data will prove valuable. You'll also find that it's worthwhile to go directly to specific staff members and get their feedback. In particular, it's often the case that a Systems Professional will underestimate the storage impact to an Application Development Professional.

At the end of the process you'll have a good enough picture to do some rough calculations. Take the total number of TB's in the shop and divide by the Full Time Equivalent staff you've just calculated. You should see figures that roughly approximate the following table:

The Impact of Staff Costs

<u>Topology</u>	<u>GB's Managed</u> (Per Person)	<u>~STAFF Costs</u> (% of total storage costs)
Distributed:	150 - 200GB	60%
Collocated:	300 - 400GB	40%
SAN of Today:	500-700GB	30%
SAN of 2000:	1200-1500GB	15%
Future SAN:	2500GB+	10%

* Source: ITcentrix 1999

This chart shows (by storage approach) the total disk gigabytes that can be managed by an individual (on average). The right column shows the percent of total storage costs (storage TCO) that is attributed to **STAFFING**.

Remember that you've just assessed an entire IT infrastructure so you'll have a mix of different types of storage (often including so-called Consolidated). Our estimate is that typical Consolidated metrics (of today) closely approximate SAN of 2000 at around 1600GB per person. The key is to focus the customer on the opportunity to improve by using a Compaq SAN approach. This means pinpointing the SAN candidates within the Collocated and Distributed environments.

Caveats:

Your customer's information will not always match these numbers

Possible explanations include:

Lack of information.

Heavy use of Clustered Systems (which approximate the benefits of SAN).

Very poor storage Capacity Utilization.

Frequently customers will demonstrate better than expected capacity management efficiency (GB's managed per person) due to very poor storage

utilization. Typically, storage capacity utilization on so-called open systems is between **40 to 50%**. In an instance where management efficiency is better than expected storage capacity utilization can drop to as low as **20%**.

Step 3: Assessing the Target SAN Infrastructure

This step uses a single screen in the tool and requires the input from five questions related to the target SAN infrastructure; including:

- # of servers in the SAN (excluding failover servers).
- Price per server (including software and maintenance). Note: The tool uses a fully loaded 4-year depreciated cost which will often range between \$20K - \$50K for NT and \$50K - \$500K (and sometimes higher) for Unix.
- Usable/Formatted disk capacity of the target SAN.
- Type of disk technology used (18GB, 36GB, etc.).
- High availability (H/A) configuration (yes or no?). Note: If “yes” is selected, the tool appropriately doubles the number of servers and attached storage to create a “spare and spare” H/A configuration.

Unless, the customer plans to install a disaster tolerant solution such as Data Replication Manager, answer No to this question.

Note: Clarify for the customer that the model assumes a usable, protected price/MB assumptions of between **\$.20 - \$.40 per MB** (depreciated and including software) depending on storage topology.



Depending on the customer, sometimes it's better to not to address this level of detail unless asked to do so. Stress with the customer that you're happy to share all assumptions associated with the model but consider leaving such well-established information such as \$/MB outside the initial assessment.

Step 4: Assessing Application Value

This step uses the input from two screens in the tool:

- Application Input
- Assessing Application Value

The Application Input Screen requires three inputs (5a, 5b and 5c). The first, 5a is simply a description of the application. The second two however, are vital for the creation of value and relate to the number of users accessing the application.

Start by asking two simple questions of the customer:

- 1.) What are the predominant applications running on these servers?
- 2.) How many total users are on these systems?

If the customer doesn't know...ask: "within a range...is it more than 10,000? Less than 10,000?" Try to narrow it down to a reasonable guess. **Typically, for Unix and NT servers, the number of users per server will range from 20-200. Use this "rule-of-thumb" to help the customer take a reasonable estimate.** If this doesn't help...narrow it down by focusing on the most important application (or applications). If this still doesn't work...take an action to follow up with someone that knows.

- 3.)What % of the users are active/concurrent?

The percentage of active/concurrent users is defined as the number of simultaneous users that are actively accessing the application at a typical point during normal business operations. The figure is meant to represent the average activity for the application during normal business hours.

This figure is crucial as it identifies those users on the system that are ACTIVELY CREATING BUSINESS VALUE BY USING IT.

The Calculating Application Value Screen provides key inputs to the ITcentrix Value Contribution Model. The process heavily relies on a proven technique used to establish the value of an application (or portfolio of applications) to the business. We use this value to calculate:

- a.) The value of better availability
- b.) The value of faster time-to-market

The basic logic of the concept is that if you have N number of active/concurrent users for an application, those users (in total) are the ones adding business value for the company. The value of that application (to the business) must equal at least the fully loaded costs of the users of that application (or else why bother installing the application).

As such, the base value of a target SAN application is established by taking the N active/concurrent users of the application times their fully loaded costs.

The slider input (6a) is used to multiply this base value by a factor (to establish the true application value) by "uplifting" the base value (just calculated). We have included some representative applications and their respective multipliers as reference but all situations will vary.

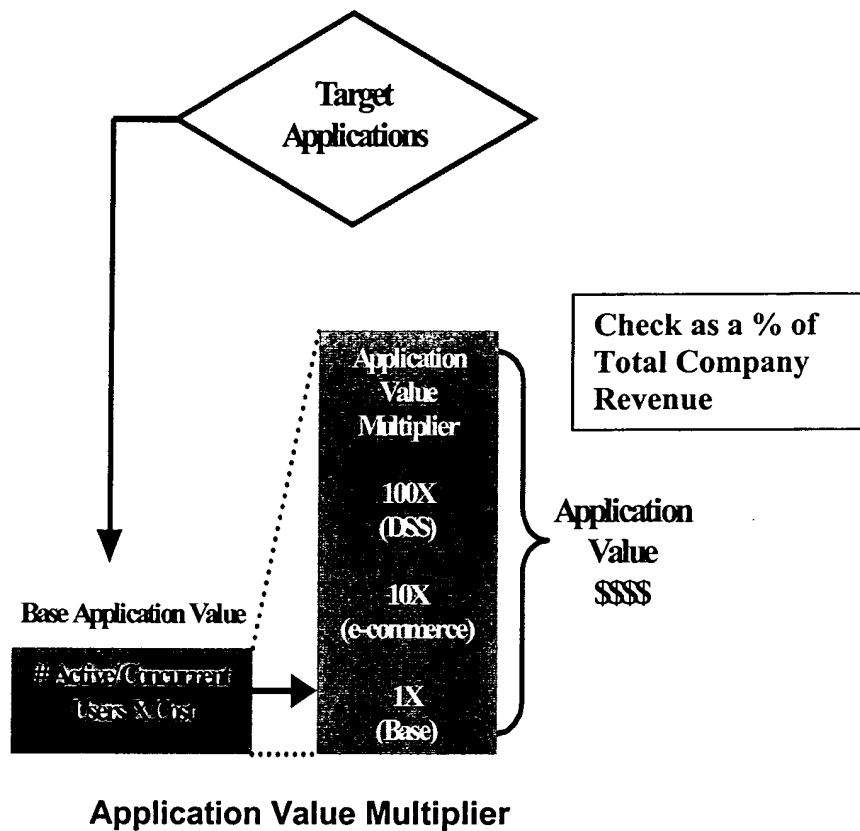
The box includes some “sanity checks” that can be used; specifically: 1) Application value as a percentage of the company’s revenue and 2) Application value per active user (which can be compared to revenue per employee).

What follows is the details behind our calculations:

To establish Application Value:

- Take the number of active/concurrent users for a specific application (or set of applications).
- Multiply this number by the fully loaded costs per active/concurrent user (we use **\$50,000** in the tool currently).
- The result establishes is the “*Base Application Value.*”
- “Uplift” the base application value by an “*Application Value Multiplier*” that reasonably reflects the total application value to the business (6a).
- Check this figure as a percentage of total company revenue.
- Check this figure on a per active user basis as compared to revenue per employee.

Calculating Application Value



EXAMPLE – Typical Customer Interaction

- Q. For these target SAN applications, what would you estimate to be the ANNUAL value of these applications to the business?
- A. That's a tough one...I really don't know.
- Well, you said there will be about 5,000 total users on the system right?
- A. Give or take...that's right.

- And we estimated that about 15-20% of those were active/concurrent users right?

A. Right.

- Ok...let's use 15% so that's about 750 active/concurrent users okay?

A. Yes.

- Now, what we do is we assume a fully loaded cost per user. What would you say that is?

A. Oh...I don't know.

- How about \$50,000 per user?

A. Sounds low.

- Okay, but let's be conservative and use \$50,000/user okay? Cause sometimes these types of users are not too expensive...

A. Right...Okay.

- Okay. So 750 users times \$50,000 = \$37.5M...Okay...are you following this?

A. Yes.

- Alright...so what we've just done is set a **Base Application Value** by taking the # of active concurrent users times their costs. In other words...the application must deliver at least the base application value to the business or else why bother? Make sense?

A. Yes, of course.

- Now what we do is begin to multiply the Base Application Value by some factor to reflect the real application value. We typically see financial applications with a relatively smaller multiplier (3X or so) and commerce applications much higher (10X or so). What do you think for these applications?

A. Well, they're pretty much back office stuff so I'd say 2-3X.

- 2-3X...okay so that's between \$75 to \$100M. Sound about right.

A. Yes, I guess so.

- Well, what's your company's annual revenue.

A. It was around \$2B last year.

- Okay, well we're saying these applications contribute roughly 5% to the company's top line...sound reasonable?
- A. Yes. 5% sounds about right...I mean they're back office but it's stuff like getting our suppliers and partners paid so it relates to the quality of our company...Sure 5%...I can buy that. That's a pretty good way to look at it...It's maybe even a bit higher than that...

Step 5: Applying Application Value to Availability

One established, it's a fairly straightforward exercise to apply Application Value to Availability. As you may recall from the description of the Availability Module we use a number of factors to assess the impact of improved availability. For the purposes of populating the model, however we, for the most part, use "default values" based on discussions with large and medium-sized customers.

For example, we assume the following assumptions about Application Availability during normal operations:

**Application Availability Assumptions
(Availability During Business Operations)**

Topology	Unix and NT Basic Configs	Unix and NT H/A Configs
Distributed	94.5%	97.3%
Collocated	94.8%	97.4%
SAN of Today	95.3%	97.6%
SAN of 2000	95.5%	97.9%
SAN of 2002	97.4%	98.7%

H/A = High Availability configuration which doubles the # of servers and storage devices.

We have also made a number of assumptions regarding factors such as scheduled downtime and % of business lost during scheduled downtime. Both have relatively minor impacts on the overall value figures.

Note: Many customers will have varying views of availability. Frequently, customers will refer to much higher availability figures but it is likely that they are not measuring availability at the application (or end user) level.

There's not much point in debating availability figures as they have a fairly small impact on the model overall. It's an important topic nonetheless and one that deserves attention. In addition, it's worth noting that the higher the application value, the greater the impact of improved availability on overall value.

The key to populating the Availability Module is answering the following question:

1.) What % of revenue (or productivity) is lost during unplanned downtime?

Low impact applications tend to see an impact of **5%** or less and high impact outages tend to **exceed 15%**. The key here is not the percentage, but the conversation that can ensue. Some sample questions follow:

Ancillary Questions:

- 2.) How mission critical are these applications?
- 3.) What's the impact of an unplanned outage?
- 4.) What do people do when the application goes down?
- 5.) What's their reaction?
- 6.) Can they transact business in other ways?
- 7.) Do they do nothing?
- 8.) Do they go home?
- 9.) What happens to the business?

Use slide bar to calculate impact of % of revenue or productivity that is currently lost during unplanned downtime for the Target SAN.

Step 6: Applying Application Value to Flexibility

As applications are enhanced through changes, business value is added. Application changes may make applications easier to use, faster, more functional, etc. As such, application changes usually mean added business value. The rate, at which changes can be made, therefore has a large impact on business value.

The tool asks the user to input the rate of application change in the environment ranging from Low to Medium to High. A low rate of change means the application is infrequently touched and has a very long useful life (e.g. more than 10 years). Payroll or many similar “back office” applications fall into this category. A high rate of change means the application undergoes significant change and has a shorter useful life (e.g. less than 5 years). Customer management applications or commerce-related activities might fall into this category.

The Business Flexibility Module uses information derived from the Application Value and calculates the value of improved business flexibility. The metric used is increased revenue or productivity from faster application development and deployment (faster time-to-market). For this version of the model, the inputs for the Business Flexibility Module are defaults based on conversations with customers. For the purposes of disclosure and discussion, the following questions can be explored:

- 1.) How frequently do you enhance existing applications?
- 2.) By roughly what % each year do you increase the number of users supported for (a specific or set of specific) this application(s).
- 3.) How many new applications do you deploy each year?
- 4.) Would having direct access to data help your applications group accelerate development?
- 5.) Would having an any-to-any topology between servers and I/O accelerate application development and deployment by simplifying space allocation and re-allocation?
- 6.) If you had direct access to data and could automatically allocate space, by what % do you think you could improve application development and deployment? 1%? 5%? 10%?

Step 7: Delivering Results

By this point hopefully you've established with the customer that you are genuinely interested in establishing business value for the customer. Some typical action items from this meeting might be:

- 1.) Follow up with other individuals to fine tune data. Examples might include:
 - A. Applications Professionals to identify storage problems
 - B. Someone who knows Staffing
 - C. An individual who has responsibility for Distributed Storage
 - D. Someone that might help clarify the Storage Capacity Utilization
 - E. A Line-of-business Professional that could help clarify Application Value
- 2.) An offer to run the Value Model and share the results with the customer.

I will provide Graphic

Outputs such as this can be pasted into proposals with other supporting documentation. The proposal could contain a Value Section that describes how a Compaq SAN delivers business value. Importantly, the models will be generated using customer-supplied data with a fairly good understanding of assumptions used in the calculations.

Reference: Appendix E Science of Selling- Letter of Understanding

Case Studies

Module 4

Purpose

What's in this Module?

The following steps are included in this module.

Topic	See Page

Appendix

Assumptions

The charts contained in the output screens have taken customer supplied information, combined with industry average data from ITcentrix surveys, and assigned values for three areas:

- Dollars from cost savings
- Dollars from higher availability
- Dollars from faster application deployment

These values are derived by overlaying key technical features of SAN (e.g. Global Naming) and assigning value contributions to them (in each of the above areas). The assignment of these values comes from detailed surveys with customers (who have SAN or equivalent technology installed). In the case of future technologies, the values are derived by probing hypothetical scenarios with customers about what impact specific features will have on their operations.

The following chart shows a summary of each function from the SAN timeline, to which topology it applies and the business impact in Cost, Availability and Flexibility terms. The model weights each of these functions and applies them accordingly to each storage topology in terms of their value contribution to the business.

Function	Topology	Cost Savings	Availability	Flexibility
Collocation	Collocated	High	Medium	Low
Sequential Volume Sharing	SAN of Today	Very High	None	Medium
Point-in-time Copy	SAN of Today	Medium	None	Medium
Alternate Pathing	SAN of Today	Low	None	Low
Fibre I/O Cards	SAN of Today	Negative	None	None
Static Volume Management	SAN of Today	Medium	None	Medium
Improved Static Volume Management	SAN of 2000	Low	None	Medium
Sequential Volume Sharing	SAN of 2000	Low	None	Medium
Fibre I/O Cards	SAN of 2000	None	Negative	None
Fibre I/O Cards with Logical Port Names	Future SAN	Very Low	Negative	None
Logical Port Names	Future SAN	Very Low	Medium	Low
Dynamic Volume Management	Future SAN	High	High	High
Dynamic Storage Allocation	Future SAN	High	None	High
Read-only access to Heterogenous Files Types	Future SAN	None	Medium	Medium
Global SAN File Management	Future SAN	High	High	High
Workload Manager	Future SAN	Medium	High	Medium
Dynamic Workload Manager	Future SAN	Medium	High	High

The following general assertions apply for each area of value:

Operational Costs:

SAN cost benefits stem primarily from the following areas:

- Significant reduction in storage management complexity.
- Reduced number of processes and procedures.
- Easier/faster problem determination and resolution.
- Automation of copy procedures.
- Simplified backup and restore.
- Significant reduction in performance management.
- Easier resolution of bottlenecks due to improved I/O pathing.
- Faster recovery process for system administrators.

Availability:

As storage topologies advance, availability improvements are primarily due to the following factors:

- Faster recovery times for hardware and software errors.
- Fewer operator errors (due to simplified IT).
- Better recovery procedures (again, simplified IT).
- Reduced backup window pressures.
- Lower probability of offline procedure (batch) overruns.
- Ability to manually balance workloads across more volumes and servers to avoid dangerously high server and storage utilization levels.
- Automated any-to-any exploitation results in significant reduction of errors related to server and storage resource constraints (Future SAN).

Business Flexibility:

As SANs are enabled, the impact on applications development and deployment can be substantial. SANs address the following storage issues:

- Applications development has better (more direct, any-to-any) access to existing data sources within the organization.
- The data for applications can be more easily (directly) extracted from existing information.
- The data required for applications can be more easily moved and exploited by other applications.
- System configuration constraints are less problematic as applications gain more dynamic access to storage and server resources.
- On balance, better data access means faster changes and simpler change management.

Glossary Terms

Business Value Contributors –

Operational Costs

- Fully loaded cost savings including factors such as staff efficiencies, hardware utilization, network costs and software costs.

Service Levels

- Incremental revenue or productivity dollars from higher availability due to reduced planned downtime, reduced unplanned downtime and improved performance (future).

Business Flexibility

- Incremental revenue or productivity dollars from faster application deployment due to reduced time to develop and introduce new applications (i.e., time-to-market) and value generated from faster deployment times.

Distributed Storage

- Physically de-centralized (outside of a data center), servers managed centrally, storage directly attached to servers and managed locally, back up over the corporate network.

Collocated Storage

- Physically centralized (inside a data center), servers managed centrally, storage directly attached to servers and managed centrally (by IT Professionals), backup over a local network within the data center.

Homogeneous SAN (*San of Today*)

- Shared, pooled storage attached to multiple Compaq servers running a homogeneous operating system using Compaq Storage. Any-to-any connectivity between servers and I/O (disk and tape) using a common serial switched fabric and network switches; static allocation of server and I/O resources.

Limited Heterogeneous SAN (*SAN of 2000*)

- Shared, pooled storage attached to multiple Compaq and non-Compaq servers running Unix (some flavors) and/or NT operating systems using Compaq Storage. Any-to-any connectivity between servers and I/O (disk and tape) using a common serial switched fabric and network switches; static allocation of server and I/O resources with improved management function.

Future SAN (San of 2002)

– Shared, pooled storage attached to multiple Compaq and non Compaq servers running multiple operating systems using multiple vendors' storage devices. Single point of control for dynamic allocation of server and I/O resources to optimize service levels and availability. Unified SAN OS or Global OS support.

Storage Area Network (SAN)

– an approach to system and storage management that enables an any-to-any fibre channel connection topology where any server (in theory_ has access to any connected storage (including tape) across at least three OS images.